

Smarter policymaking through improved collective cognition?

Anders Sandberg, Future of Humanity Institute, Oxford Martin School, Oxford University

Preprint of chapter in *Anticipating 2025: A guide to the radical scenarios that lie ahead, whether or not we're ready for them*. Ed. David W. Wood. London Futurists. 2014

Policymaking is far more than forming or expressing opinions. Having a voice matters, but so does setting values, spreading knowledge, considering solutions, negotiating agreements, taking responsibility and checking what is going on. New technologies may help us innovate smarter ways of doing this by exploiting the collective intelligence of humans.

What is policymaking, really?

Governance refers to "all processes of governing, whether undertaken by a government, market or network, whether over a family, tribe, formal or informal organization or territory and whether through laws, norms, power or language." (Bevir 2013) It is not just a matter for governments: we are doing it every day. The governance done by states and organisations just happens to be more formal. It can also be more or less fair (how much the 'deciders' respect the rights and interests of the stakeholders), efficient (how much effort decisions require), legitimate (whether the stakeholders agree that the deciders should decide), and of course more or less successful at achieving its goals.

Governance typically involves policies: a statement of intent, and an implementation of this as a procedure or protocol which will then guide what will be done. Policymaking is deciding on the policies.

A common model for this process is the policy cycle:

- **Agenda setting:** "Somebody needs to do something!" (Identifying that there is a problem that demands attention.)
- **Policy Formulation:** "What options do we have? (and who are we?)" (exploring options available for addressing the problem. This may include determining who the stakeholders are and getting them involved.)
- **Decision-making:** "This is what we will do!" (The executive determines what the course of action will be.)
- **Implementation:** "I am doing it." (Putting the decision into practice.)

- **Evaluation:** “Did it work?” (Assesses the effectiveness of the policy in terms of the intentions and results; new agenda-setting may be required.)

This is a convenient oversimplification since in the real world the steps often overlap and mix in complex, interactive ways. Map and policy reality can be very different (Hallsworth et al. 2011).

An important point is that the system should be open. As noted by David Brin (1998), open societies can correct their own errors. Members point out what they consider to be problems, and if enough agree the issues get handled – people are held accountable, failed policies changed, new institutions created. Closed societies are only self-correcting if the elite chooses to observe and correct a problem. This tends to prevent corrections dealing with malfunctions in how the elite run things. Over time inefficiency, corruption and risk builds up, leading to disaster.

Challenges ahead

Governance is facing interesting challenges. It has to deal with a more diverse population where one cannot assume everybody will have roughly the same preferences or even ideas about what is going on. The world is becoming more globalized, meaning that actions – both individual and governmental - will have transnational effects. We are living longer, making the electorate consist of several generations with increasingly different life experiences, values and preferred technologies. The problems facing the world are increasingly complex, and some priorities such as global catastrophic threats may be becoming more extreme.

Technological change can act to change the policy rationales (Foldvary & Klein 2003). As transaction costs are reduced, many market failures disappear. Regulated taxi drivers might be necessary if there is no easy way of guaranteeing driver and customer safety and that payments are properly made and taxed, but the introduction of app-based transportation network companies (TNCs) such as Uber and Lyft resolves many of these problems and casts doubt on existing policies. Often technological changes lead to old policies becoming roadblocks, sometimes maintained by incumbent interests (as witnessed by the fierce resistance against the TNCs in many places). Conversely, technology may increase the connectedness of systems, producing new systemic risks that are by their nature transnational, fast-moving and outside the experience of current governance structures.

Bad governance *matters*. In the World Economic Forum’s risk report the “global governance failure” risk is central: when good decisions cannot be made, every problem will get worse (WEF 2014). The most repressive governments are typically also the worst in terms of human development indices, and even without ill will bad policymaking means that alleviating poverty will be hard.

Looking at what events have killed the most people in history the main threats appear to be pandemics and the effects of bad governance: wars and genocides. As our technologies of destruction and surveillance become more powerful the need to keep states benign becomes more urgent. This requires better governance.

Where have we been and where are we going?

(Moore & Hartley 2008) looked at modern governance innovation and pointed out that unlike how innovation in business is typically described, governance innovation is not just better products and processes. It can involve changing the organisational boundaries, tapping new pools of resources, exploiting government capacity to shape private rights and responsibilities, and redistributing the right to define and judge value.

For example, governments can outsource some of their activities to companies, community groups or individuals (or conversely, take over previously non-government roles). In fact, the boundaries might stretch into distributed networks inside and outside the government where decision-making and optimization happens everywhere. This can tap new pools of resources: not just money, but information, human skill and energy. One of the great successes of social democracy was the ability to exploit the untapped resources of previously disenfranchised people: not just for protest or as an influx of productive members of society, but to build institutions like co-ops, unions and popular education. On the more sinister side, systems of citizen informers - from the Venetian "Bocca di Leone" anonymous denunciation system, to the use of 2.5% of the East German population as informants by the Stasi – also allowed the use of local, normally untapped information to reach the authorities.

Technology and governance innovation

Obviously technology is an important enabler of governance innovation. With better communications interactions with the government do not have to be face to face (especially if there is a way of verifying that the parties are who they claim to be). Congestion charges (incentivising individuals to meet social objectives) are possible thanks to computer vision and vehicle databases. If activities can be automatically documented, accountability can be maintained even when outsourced.

Political systems are information limited: the kinds of information storage, transmission and processing that can be done determine what political systems are possible.

Without writing there cannot be a code of law. Certainly there can be traditions recited by memory (employing the technologies of rhyme and memorisation), but their complexity and reliability will be limited. With writing uniform laws could be developed. The Roman and medieval introduction of punctuation systems, separation of words, paragraphs, tables of contents and Arabian numerals led to modular texts where parts could be handled as individual chunks, allowing amendment processes rather than radical rewriting: now rules could be updated piecemeal.

Without printing parliamentary democracy cannot work, since party platforms need to be distributed to the voters, who need to be sufficiently literate. Mass democracy requires broadcasting (either through newspapers or radio), and so on. We have recently gained a huge increase in information management through computers and the Internet but we have not yet explored much of the space of possible new forms of governance.

History shows that innovation in governance requires more than just information processing. The printing press arrived in the West in the 15th century and certainly contributed to dramatic political change. However, governance innovations continued: mirrors for princes (instruction books for young rulers; implied that ruling was a learnable skill), the elaborate Venetian election system

protocols (discovering how to design divisions of power hard to cheat or hi-jack), parliaments, the elections of NCOs in Landsknechte companies, reforms of the feudal systems, the emergence of the Westphalian system of sovereign nation-states, permanent civil services, collegiate bodies, drafting organisations, committees, regularized meeting rules, formal government bookkeeping and statistics, just to name a few. By the 19th century the national state was run by a professionalized civil servant class with meritocratic ideals, loyal to the state.

Printing enabled or empowered much of the above but the structures also had to grow organically. Societies run on trust and social capital. The rule of law is not just about the existence of law and law-enforcement, but also the internalisation of the habits of following the law and expecting the law to be followed. Trusting formal structures allows far lower friction in running a society than relying on personal contacts, but one needs to grow up in the system to understand and trust it.

We should hence not expect that the ideas for new policymaking suggested by recent technologies are the final word: in fact, they are merely the earliest beginning. Even with far more rapid social change than in the past it will take years or decades before we invent and adapt to the governance innovations that will come out of the information revolution.




Figure 1: Google n-gram search for six communications technologies 1800-2000, showing in what fraction of texts they occur. The Internet shows a remarkably fast and strong rise in our cultural consciousness.

Group problem solving

Group problem solving lies at the core of policymaking: several people need to work together to implement the policy cycle, either because they are stakeholders or just because there is much work to be done.

Simple organisations have a “logistic bottleneck”. Each individual spends their time doing useful work. If all work can be done in parallel the total amount of work will grow proportionally to the number of people. However, in real situations some communication is needed between the members: they have to spend some of their time talking rather than working. If each member needs to talk to everybody else, the talk time grows proportional to the size of the organisation. This means that the total amount of work done will peak at some size. Beyond that point, each new member causes greater communication losses than they contribute useful work (Figure 2).




Figure 2: Graph of useful work for a simple organisation as function of the number of people in it, assuming each worker needs to spend about 3 minutes coordinating with each other worker and works 8 hours per day.

This sets a limit to the size of the organisation. One solution is to split it into a hierarchy of smaller parts that mainly communicate internally. This enables larger organisations and more work to be done, but only functions if the tasks to be done fit the information flow. It also places demands on the managers to spread the relevant information horizontally: they can easily get overloaded or miss relaying important news. Also, as the number of layers of organisation grows, the probability that messages going up or down the hierarchy get distorted increase; large organisations are also limited by the need for control (Williamson 1967).

Information technology clearly can improve things. Used *right* it can help reduce information overload. Faster or more selective communication methods reduce communication time losses. If messages are less likely to be lost or distorted cohesion improves. Cheap long-distance communications allow greater geographical dispersal. Storable communication (letters rather than face-to-face communication, emails rather than phone calls) allows temporal dispersal: not everybody needs to be working at the same time, and parts of the organisation might act at different speed.

Research in group problem solving has shown that it can be good for smoothing out individual bias (the famous “wisdom of crowds” effect), either because the median or mean is likely close to the true answer (Galton 1907, Condorcet’s Jury Theorem), or because internal processes lead to convergence (market pricing, the scientific community). Diverse groups can approach problems from multiple angles and ideally combine distributed knowledge that no single member possesses. In particular, groups can be better than their best member at solving problems, whether deception detection (Frank et al. 2004), solving word problems (Laughlin et al. 2006), recall (Smith et al. 2009) or collaborative reasoning (Moshman & Geil 1998). This is especially true for problems that have solutions whose correctness can be checked.

Unfortunately, performance is far worse on “wicked problems” where there is no agreement about their definitions, much less their solutions or even by what values to judge them (Sandia 2007, Stevens et al. 2008). This is particularly problematic from a policymaking standpoint, since wicked problems are just those we most wish to solve (Head 2006). Groups can also fall prey to groupthink

(members refrain from voicing criticism out of worry about their position), polarization (especially when some members have strong and unyielding views), social influence (Lorenz et al. 2011) and numerous other cognitive biases (Sunstein 2006). Add to this the above scaling issues as the group becomes large, and it is clear that there is much room for improvement in group decision making (Kerr & Tindale 2004).

The potential stupidity of groups is a real problem in governance and policymaking. The “madness of crowds” has been a recurring worry in governance since political philosophy was first invented, often being seen as an argument against democracy. But it applies to any multi-person government too. Most big decisions are not made individually but as a result of group decisions: history is littered with examples of how disaster followed (the Bay of Pigs Invasion, WW I, financial bubbles etc.) There is a negative correlation between the Human Development Index of countries and the number of members in the cabinet (Klimek, Hanel & Turner 2009). Even fairly well-functioning governments continually make disastrous blunders due to failures of group cognition (King & Crewe 2013). Improving group cognition hence seems to be a key challenge for handling a more complex world.

New possibilities

Network media

The most obvious change over the past few decades has been the democratization of communication: from a world where one-to-many communication was expensive and often monopolized/controlled by the state to a world where many-to-many communication is cheap, ubiquitous and hard to control. This is manifestly important for opinion formation and distribution.

But the free and flexible structure of the new media has other important effects. It allows the formation of new networks, including discovering other people with shared causes. This disproportionately helps people with rarer views to find each other and form groups, increasing the diversity and inhomogeneity of the social space. The ability to select information sources also allows the emergence of informal information channels that efficiently channel relevant information to the people who want it. This enables rapid, synchronized responses without a central coordinator.




Figure 3: Simple model of network media self organisation. (Left) The initial random network can only weakly transmit information of interest between agents. (Centre) As the links are updated, the network exhibits larger information cascades distributing information to agents interested in it. (Right) The growth of the average information cascade over time.

A simple model can demonstrate this phenomenon (Figure 3). Agents (bloggers, twitter users, etc.) in a network read a limited number of information sources (other bloggers and twitter users) and have interests in certain topics. Random agents post original information on topics based on their interests. If one source has a piece of interesting information that source is regarded as more useful by the recipient agent, who might also write about the information. This can trigger information cascades as their readers in turn react, spreading the information in the network. Over time agents turn to other (randomly selected) sources if one source proves useless.

While simple, this model produces a collective learning behaviour. From the start agents with similar interests are unlikely to be linked. News rarely spreads far. But gradually topic-specific sub-networks emerge that transmit information to all interested agents. This happens despite each agent having limited input channels: they cannot listen to everybody, yet they find an organisation that allows collective sharing using only local information and efficiently shunts it to most people.

Social media also have the effect of enabling people to find “microexpertise” and “microteaching”: informal requests for explanation, discussion and help to friends-of-a-friend within the network. These local channels of expertise are more trusted since they are embedded in people’s social network.

Different media enables different forms of expression and hence have different policy effects. Twitter is instant and ‘pushes’ but has low discursive bandwidth, suitable for alerting about something. Blogs enable semi-permanent expression and much deeper arguments but are ‘pull’ media that need to be sought out, with the potential for interactive discourse with interested people. Online forums focus more on the discussion, and can be selective for particular communities.

It is worth noting that thanks to the information revolution new media are likely to be innovated at a high rate. In the past, new media such as print or radio required an entirely new technology and building a new infrastructure. Modern media piggyback on an existing infrastructure, allowing a very low threshold for entry. If somebody has an idea it can usually be implemented cheaply, democratizing media innovation, not just content innovation. Consequently we should expect much experimentation with new forms of media, some of which will no doubt affect policymaking. While it is nearly impossible to say anything about their impact, it seems plausible that some of them will weaken the power of existing media – and hence institutions based on them. The broadcasting- and newspaper-based 20th century model of how public discourse is shaped is likely doomed.

The noise cycle

There is also a dark side to networked media. As information moves between people it is often selected and distorted, even if the underlying media allows perfect copying. Traditional media act as a single level feedforward network: media consumers can select who to listen to, finding sources they agree with (if there are any). Network media are multilevel: information gets copied, amplified and modified by several intermediaries. The challenge with network media is that the extra layers amplify the noise and bias in the transmitted information like the “Chinese whispers” game.

Even if the information itself remains unchanged, whether it is transmitted or not may depend on local bias. For example, consider two scientific studies honestly reaching opposite conclusions about the safety of something common. The one suggesting more risk is more likely to make an eye-

catching press release, improving the probability that a news outlet will react to the press release. Readers are more likely to look at a news story suggesting they are at risk than one suggesting no risk. If the readers then respond by blogging or tweeting about the story, the same biasing process repeats, possibly several times. Here the bias of risk salience make one study widely reported, while the other, disconfirming study is underreported or not even reported at all. If there are other cultural or biasing factors (such as mistrust of technology, an ideological commitment to certain views, or just scepticism of hearing disconfirming stories) the effect can be magnified even more. Not to mention the effect of accidental or deliberate distortion besides selective reporting.

There is often a speed-accuracy trade-off in collective decision-making. Often a headline and a summary gives enough surface information to decide on how one will likely react to a piece of news: many hence respond to that image rather than reading the full story or going back to the original source. It is easy and quick to respond the way one's neighbours do, conveniently assuming they have done more analysis than oneself. Ethologists have observed the same issues in social animal behaviour, where information cascades lead to arbitrary choices (Giraldeau et al. 2002). Incorporating the right kind of negative feedback to remove noise, to reintroduce original unbiased information, and to criticize facile decisions, is necessary both for human and animal communities that want to reduce their sensitivity to individual errors and detect the real patterns in their noisy environment (Couzin 2008). From a policymaking perspective this is where new technologies may now be needed the most in order to turn the network media into a force for better policymaking rather than worse.

Reputations

One way of handling the distortion risk and problems with disruptive people is reputations. In society reputations are the public beliefs held about us, strongly affecting how we are treated by others. Reputations are constructed by social networks so they are somewhat robust against individual dissent and gaming, and since they affect treatment people are motivated to maintain good reputations. They help stabilize social structures, make trust go a bit further, and can reduce transaction friction – something very useful for network media and crowdsourcing (Botsman 2012, Swallow 2012).

In many online situations they are little more than scalar values (upvotes, karma, ratings etc.) that have the benefit of being easy to quickly check, but as interactions become repeated certain online identities gain true reputations. The value of having a good reputation both makes people more likely to commit to a community (since they lose the advantage if they leave) and to behave well inside it. "Trust tokens" are useful for constructing online teams: getting badges, referrals or signals about social capital (such as number of friends or reviews) help people build trust more quickly (Morita & Burns 2013).

However, reputation systems are imperfect in many ways. Merit and reputations are not necessarily strongly correlated. Thanks to rich-get-richer phenomena (the 'Matthew effect') well-known people are likely to become more well-known: you can become famous for being a celebrity. This means that the merit part of a reputation can be overshadowed by the celebrity part. Well cited scientific papers are likely not *bad*, but quite possibly well cited because they were easy to find.

Reputations are also becoming more global and more permanent thanks to new media. A damaged reputation may hence not go away if one moves to Tasmania: Google, Facebook and Wikipedia will

be just as present over there. Attempts at legally protecting reputations by for example censoring online information (as a number of European court cases have attempted) typically fail because of the “Streisand effect”: it draws attention to the issue, almost guaranteeing a stronger effect. Reputation management is widely used, and of course involves attempts at burnishing one’s own reputation regardless of what others actually think or what one has done. The equilibrium is however likely more on the side of truth than spin: the rough truthfulness of reputations is why they matter. If spin was too easy reputations would not matter. However, in domains where reputation management is so costly that only a few can do it but reputations otherwise remain truthful, there is a risk that elite reputations become misleading.

A world of global, easily checked reputations means that politicians and other people involved in the policy process will be scrutinized to a degree never seen before. Conversely, it can also help stabilize online collaboration and crowdsourced policymaking.

Crowdsourcing

Crowdsourcing is the current banner for new methods of getting things done by getting partial contributions from large volunteer groups of people, who typically benefit from it (Howe 2006). The difference from previous collaboration methods is that information technology lowers the cost of getting large numbers to contribute, reduces the coordination costs, helps shape the interactions in useful directions, and may allow automatic joining and quality control of the “workpieces”.

Crowdsourced research (“citizen science”) has a long history (amateur astronomers and birdwatchers reporting sightings, genealogy research) but truly began to take off when connected with web platforms. The good experiences with the “Galaxy Zoo” (www.galaxyzoo.org) led the Oxford University astronomy department to develop a platform (the “Zooniverse”, www.zooniverse.org) for citizen science, which in turn has led to methods of combining individual contributions in statistically optimal ways (Simpson et al., 2013). Cloud labour systems like the Mechanical Turk allows human abilities that are hard to automate, such as pattern-recognition or common sense knowledge, to be “put on tap”. Wikipedia and open source software demonstrates how very complex projects of collective knowledge and open innovation can be channelled and self-organize with the right software support frameworks. Under the right circumstances this can produce very advanced collective intelligence when working on open-ended puzzles (McGonigal 2008). More recently methods for crowdfunding such as Kickstarter have allowed new methods of funding, especially for niche products or ideas.

The real skill required for successful crowdsourcing is to design the system so that it generates the right kind of interaction that boosts productivity while filtering away errors or low-quality signals, as well as manages to attract a critical mass of contributors. This is a discipline as yet in infancy. Given the tremendous potential in crowdsourcing there are strong incentives – economic, social, scientific - for people to try to develop this discipline. Once crowdsourcing innovations happen they can typically easily be copied and modified. Hence I am confident that over the next decades we will see dramatic improvements in the ability to *reliably* design crowdsourcing systems for particular purposes.

Crowdsourcing can obviously be applied to policymaking. (Brabham 2013) suggests four main approaches to crowdsourcing in government: (1) knowledge discovery and management (collecting

distributed knowledge), (2) distributed human intelligence tasking (distributing micro-tasks that require human intelligence to solve), (3) broadcast search (pose a problem solving challenge, allow the online community to find a solution), (4) peer-vetted creative production (the community both proposes solutions and collectively chooses among them).

Petition tools like change.org and avaaz.org serve to influence decision-makers by focusing stakeholder preferences, showing there exists an opinion for or against an issue. While one can argue that this is merely a form of “slacktivism”, demonstrating preferences without doing anything, from a decision-making perspective knowing *about* the preferences might be relevant – policymakers are regularly surprised when they discover there exist large constituencies caring about what appears to be neutral issues. However, due to the multiplicative effects of network media this might not be an accurate measure of actual stakeholder views. Other approaches such as sentiment analysis (which looks at positive or negative associations to concepts in social media use) may provide a different source of preference information. Prediction markets such as the Iowa Electronic Market can be used to estimate views and likelihoods in politics. Here individuals are motivated by potential economic gain to generate accurate predictions, with the market mechanism integrating their individual guesses.

Other approaches involve civic engagement systems, whether acting on a local, national or global scale. Using apps to report problems such as potholes (the canonical example) allows distributed real-time information gathering, but is still a one-way tool. The true importance of citizen engagement is when it gets people *involved* in the activities of civil society: this allows better understanding of social needs, more diverse voices and ideas, boosts legitimacy of projects and decisions, and allows “buy in” into responses to challenges (Simon & Davies 2013).

Crowdsourcing is a very clear example of how technology can allow governance innovation in all of the areas mentioned by (Moore & Hartley 2008): modifying organisational boundaries, tapping new pools of resources, changing rights and responsibilities, and distributing the right to define and judge value. The real challenge ahead is the latter two categories: while most governance systems are fine with getting better information, problem solving and resources, they often resist influence coming from outside their boundaries. But crowdsourcing by its nature will tend to redefine those boundaries even if the incumbent system does not. While it is possible to use crowdsourcing without allowing it to change how governance works, the experience so far with successful client-crowd interaction suggests that it benefits from intimate interaction and the formation of a shared community – that way the governance (both original and over the crowdsourcing community) also gains legitimacy.

Where will we be in 2025?

Better agenda setting

Network media allows many more people to participate in agenda setting, both in noticing problems and to announce their support for resolving them. The increased diversity of communities means a better chance of someone noticing issues, but also far more competition for setting the agenda.

Deliberate knowledge discovery where the policymakers poll communities for agenda issues seems doable, but since the bottleneck is the later stages of the policy cycle it is likely that there will

already be more than enough spontaneously suggested agendas. However, deliberate searches for underrepresented voices might be important both for diversity and legitimacy. One of the side effects of bottom-up agenda setting is that groups that are more active or have more social activist capital will get their agendas overrepresented (the canonical example is again the bias introduced by pothole detection apps – they make city governments repair them in relatively affluent areas where many people use smartphones). Here top-down agenda search may play an equalizing role.

Better policy formulation

Policy formulation is where broadcast search and peer-vetted creative production come into their own. In fact, traditional requests for public comment or departmental referrals can be seen as precursors. Network media enable broader and more diverse groups to generate far more policy options. Another benefit is that crowdsourcing policies may reduce cultural disconnect and group-think (major reasons for government blunders, as outlined in (King & Crewe 2013)) among policymakers simply by making them interact more with citizens.

Some of the new possibilities that have been tentatively explored involve citizen participation in commenting on proposed bills and existing laws such as the Madison project (Finley 2013), and the Icelandic constitutional reform 2010–13. “Wikidemocracy” where citizens comment on or even propose bills within a parliamentary system does not seem far-fetched, although it has yet to find its form or be applied.

Several software initiatives for improving public deliberation and participation exist, such as loomio.org, democracyos.org and liquidfeedback.org. In many cases the software involves experimental political ideas or practical tradeoffs: DemocracyOS does not allow anonymous participation (hence strengthening reputations, accountability and transparency) and ballots are not secret (Finley 2014). Liquidfeedback also avoids secret ballots, implements voting delegation and preference voting, and has forms of structured deliberation designed with minorities in mind (Behrens et al. 2014). Software tools can implement fairly complex proposal display and voting mechanisms that would normally be hard to implement in traditional forums.

While the above systems aim at encompassing several steps of the policy cycle, other tools aim at improving the deliberation step. Decision support systems linked to open databases and simulation tools can allow exploration of options. Argument map software attempts to make it easier to trace the arguments and rebuttals in an issue. There have been attempts to implement deliberative democracy and deliberative polling online in order to foster better engagement and less polarized views (Carpini et al. 2004, Luskin et al. 2004).

Such platforms will never involve *all* citizens even if they could scale to that size because of “rational ignorance”: most people tend to leave politics (as long as it runs well enough) to interested parties. More problematic is that platforms depend on having a sufficiently large community in order to function and have legitimacy. Since online communities and platforms need to be sought out (they are “pull media”) their internal function requires an ongoing critical mass of members generating new discourse. Should that ever fall below a critical level the community goes silent and the value of participating decreases, reinforcing the decay. Joining incurs the cost of learning the system and becoming part of the online community: gaining a reputation and learning the internal norms is a nontrivial social task. So there is both a bootstrap problem in starting a new community and a

maintenance problem in ensuring that it remains active enough to be viable (and avoids sclerosis as highly invested old members seek to maintain the status quo). These problems are universal for online communities and forums: while we do not have reliable design criteria for handling them yet, it seems likely that by 2025 they will be far better understood and “designer communities” will be common.

Better decision-making

Decision-making is in theory just selecting the best policy option from a set. Unfortunately the measurement of how good the options are can be problematic, the goodness may be relative to an evaluation scale different stakeholders disagree on (requiring negotiation and compromise), and in practice many people select actions by taking the first action that pops into mind (primed by the current situational context) and seems good enough (Klein 1999).

That real-world decision-making might be more about template matching than expectation maximization is not necessarily a problem, as long as decision-makers can tell whether the situations actually fit their experience and have reasonable evaluations. Some emerging tools may help improve the pre-decision evaluation step. Simulations allow testing out consequences of some policies. Data mining, methods for unobtrusive A/B-testing, randomized controlled trials (Haynes 2012) and rigorous evidence-based policy (Head 2006) may give better evaluations. Combining expert knowledge into more reliable predictions or estimates can be done using software or deliberately structured groups (Mellers et al. 2014). The lacking piece may be training for non-expert decision-makers and better tools for negotiation among non-experts.

While the new large diverse stakeholder groups are positive for policy formulation they pose a big challenge to the legitimacy of traditional decision-making (since they are hard to fully represent, and might extend outside the remit of the traditional decision-maker, for example internationally).

Another worry about governance in a globalized diverse world is that many governance mechanisms appear likely to deadlock when too many groups are involved. One approach that has been proposed is to abandon the idea of a United Nations-style universal governance system and instead aim at “coalitions of the working”, groups that voluntarily form multilateral teams to solve joint problems and welcome buy-in from other groups (Oxford Martin School 2013). This has many similarities to IETF and other loosely networked organizations, without assuming full adhocism. It might be that the same approach could be applied for local governance, implying the formation of dynamical networks of varying levels of formality to solve problems locally.

Distributed decision-making via crowdsourcing also makes tracking responsibility for decisions harder: if a community misbehaves, how can blame be allocated? It might be possible to use the documentation of the deliberation to track accountability, although there is a risk in attempting to formalize something that is a somewhat fuzzy human concept into code: key elements might be lost. Worse, the new nature of the method can lead to *accidental* extreme concentrations of responsibility that are very hard to predict. A post on a forum can be read by hundreds of thousands, affecting their behaviour. If a fraction of the readers riot, how much responsibility should be assigned to the original poster? Given the unpredictability of which posting will take off virally it seems unfair to assume a high degree of responsibility for any individual post. The same may be true

for crowdsourced policymaking, but would imply a very low degree of accountability for crowdsourced policies.

Manual vote counting represents an interesting example of distributed human tasking: the basic operation can largely be automated, but the procedure serves to establish trust and a sense of community. It would be interesting to study whether similar positive effects could be achieved in online communities by shared tasking where potentially opposed stakeholder groups work together in maintaining the procedure and the agreed overall system.

Better implementation

One increasingly common way of improving implementation is to outsource it. Crowdsourcing suggests ways some implementation tasks might be outsourced not just to for-profit or non-profit organisations but also to volunteer or paid groups coordinated online. If they have been involved in the earlier policy formulation step buy-in and understanding will likely be stronger.

Outsourcing could in principle do 100% of implementation. Robin Hanson (2000) has suggested a system (“futarchy”) that outsources implementation completely, leaving policymakers to set values and goals, and an information market acting as selector of which implementation offer is best.

In software systems the enforcement of rules is often done automatically through code; in the future the pervasiveness of code and automation may enable automatic policy implementation. This would enable rapid implementation, use less manpower and self-documenting implementations. However, automated policy implementation will not automatically be fair, transparent, effective, or fit for purpose. It is easy to find examples in current software where discrimination is accidentally built in, information about usage is not available, and the end result falls far shorter than expectations. Designing complex systems is fundamentally hard.

In some cases like cryptocurrencies such as bitcoin the outsourcing has been deliberately designed to preclude certain forms of governance (in this case central banking and tracking of transactions). Bitcoin represents an institution created to act as an alternative or successor to current institutions, possibly with the potential to act as a platform for other functions such as self-enforcing contracts (vbuterin 2014). It is worth recognizing that sometimes policymaking consists of “exit” rather than “voice” in a community or institution (Hirschman 1970), and technology may make some forms of exit easier. This can be fundamentally destabilizing, but a world of institutions competing for participants may also be a place where they have a stronger incentive to provide constituents with value.

Better evaluation

This might be a domain where distributed human tasking is particularly effective: stakeholders distributed in the field can report back how the policy affects them or even perform micro-investigations.

The evidence based policy movement has promoted rigorous testing of policy outcomes; while practical issues, value differences and unwillingness to put one’s ideology on line may preclude testing of many key policies, there is clearly much potential for improving policy evaluation. One challenge that is worth tackling is giving stakeholders tools to investigate policy outcomes, including possibly the ability to suggest and implement policy experiments.

While Big Data may be getting the headlines, the real change occurring right now may be the emergence of Big Input. Getting the data into the system has until recently been the major problem (entering data by hand, scanning documents, uploading video streams), but when the default mode of objects in the Internet of Things is to record data using their onboard sensors the world essentially becomes one big sensor network. The growth of massive surveillance systems and identity technology strengthens this trend. From a policy evaluation perspective this is manna from heaven, assuming the data can be accessed and analysed.

The growth in surveillance capability seems largely unstoppable, but its consequences depend sensitively on how much surveillance – Little Brothers looking back – can be brought to bear on Big Brother to enforce accountability (Brin 1998). An interesting demonstration is police monitoring in Rialto, California where the introduction of video recording of police officers has led to a large reduction in the use of force – and complaints against the police. When the systems involved in policy implementation (whether law enforcement or procurement) can be automatically audited, incentives for accountable behaviour improve.

Debugging policymaking

Detecting and correcting problems with policymaking is itself a policymaking problem. Real governance systems continually struggle with it (Hallsworth et al. 2011, King & Crewe 2013).

One of the problems with improving policymaking is that we typically cannot see rare problems unless we test longer than the MTBF (Mean Time Between Failures). A problem that typically occurs once per year will not be observable if we only test for a month. In software the solution is to use multiple beta-testers: if twelve people use the software for a month, they have a good chance of catching the once-a-year problem (assuming it is not seasonal). The more people who test software, the more bugs can be found. Unfortunately it is hard to parallelize testing for society-scale innovations: there is only one political system, and bugs may only show up every few years no matter how many people use it simultaneously (but at that point, the bug may of course affect them all). Insofar it is possible, it is hence a good idea to try out governance innovations on several separate testbeds, ideally on different scales: that way problems can be discovered before the full system is rolled out nation-wide.

Policymaking supported by crowdsourced collective cognition needs to deal with signal-to-noise ratios in communication far more than policymaking done by a small group of vetted or self-selected policymakers. Negative feedback to counteract biased information cascades in network media, reputation management to encourage useful feedback and discourage malicious users, and handling large amounts of noisy data in a robust manner will be required.

In a world of fast communication and technological change, determining the appropriate timescale of policy change becomes important. It is commonly observed that constitutions are intentionally made harder to change than laws, which in turn change more slowly than regulations, strategies and projects. This way core assumptions and rules that underlie the more changeable layers remain stable, secure from brief fashions or moral outrages, without precluding adaptation to new circumstances. However, the speed at which policies must change is also itself potentially changeable. Detecting that policies are lagging and then updating the change procedures is an

important part of the meta-agenda (as is detecting that policies might be fluctuating too wildly and need slowing).

Closing

We cannot predict how policymaking in 2025 will be done. But we can see areas where innovations would improve the chance of good outcomes: better ways of incorporating de-biasing in network media and group decision-making, methods of opening up government data for public inspection, automating accountability, tools for combining preferences, tools supporting negotiation, and frameworks helping non-experts to formulate viable policy proposals.

By 2025 we will no longer notice the collective intelligence of network media embedding us, just as we today no longer notice the computers that litter our environment. With some luck we might however notice that policymaking has become a bit clearer, cleaner and cleverer.

References

Behrens, J., Kistner, A., Nitsche, A., Swierczek, B. (2014). *The Principles of LiquidFeedback 2014*.

Available at: <http://principles.liquidfeedback.org/>

Bevir, M. (2013). *Governance: A very short introduction*. Oxford, UK: Oxford University Press.

Brabham, D.C. (2013). *Using Crowdsourcing in Government*. IBM Center for The Business of Government. Available at:

<http://www.businessofgovernment.org/sites/default/files/Using%20Crowdsourcing%20In%20Government.pdf>

Botsman, R. (2012) Welcome to the new reputation economy. *Wired* September 2012.

Brin, D. (1998). *The transparent society*. Perseus Books.

Carpini, M.X.D., Cook, F.L., Jacobs, L.R. (2004). Public deliberation, discursive participation, and citizen engagement: a review of the empirical literature. *Annu. Rev. Polit. Sci.* 2004. 7:315–44

Couzin, I.D. (2008) Collective cognition in animal groups. *Trends in Cognitive Sciences* Vol.13 No.1. 2008. pp. 36-43.

Klint, F. (2013). Out in the Open: Hackers Bring Lawmaking Into the 21st Century. *Wired* 10.21.13.

Klint, F. (2014). Out in the Open: An Open Source Website That Gives Voters a Platform to Influence Politicians. *Wired* 05.05.14.

Foldvary, F.E., Klein, D.B., eds. (2003) *The Half-Life of Policy Rationales: How New Technology Affects Old Policy Issues*. New York University Press, 2003.

Frank, M.G., Paolantonio, N., Feeley, T.H., Servoss, T.J. (2004). Individual and small group accuracy in judging truthful and deceptive communication. *Group Decision and Negotiation* 13(1), 45–59.

Galton, F. (1907). Vox populi. *Nature*, 75, pp. 450–451.

Giraldeau, L-A., Valone, T.J., Templeton, J.J. (2002) Potential disadvantages of using socially acquired information. *Phil. Trans. R. Soc. Lond. B* 2002 357, doi: 10.1098/rstb.2002.1065

Hallsworth, M., Parker, S., Rutter J. (2011). *Policy making in the real world: evidence and analysis*. Institute for Government.

<http://www.instituteforgovernment.org.uk/sites/default/files/publications/Policy%20making%20in%20the%20real%20world.pdf>

Hanson, R. (2000). Futarchy: Vote values, but bet beliefs. Retrieved 2014-05-07.

<http://hanson.gmu.edu/futarchy.html>

Haynes, L., Service, O., Goldacre, B., Torgerson, D. (2012). *Test, Learn, Adapt: Developing Public Policy with Randomised Controlled Trials*. Cabinet Office and Behavioural Insights Team. 14 June 2012.

Head, B. (2009). Evidence-based policy: principles and requirements. In *Strengthening Evidence-based Policy in the Australian Federation - Roundtable proceedings - Volume 1: Proceedings*. Australian Government Productivity Commission. pp. 13-26.

Hirschman, A.O. (1970). *Exit, Voice, and Loyalty: Responses to Decline in Firms, Organizations, and States*. Cambridge, MA: Harvard University Press.

Howe, J. (2006). The Rise of Crowdsourcing. *Wired*. 14.06 2006

Kerr, N.L., Tindale, R.S. (2004). Group performance and decision making. *Annual Review of Psychology* 55, pp. 623–655.

King, A., Crewe, I. (2013). *The Blunders of our Governments*. Oneworld.

Klein, G. (1999). *Sources of Power: How People Make Decisions* Cambridge, MA: MIT Press.

Klimek, P., Hanel, R., Thurner, S. (2009). Parkinson's Law quantified: three investigations on bureaucratic inefficiency, *J. Stat. Mech.* P03008 (2009) doi:10.1088/1742-5468/2009/03/P03008

Laughlin, P.R., Harch, E.C., Silver, J.S., Boh, L. (2006). Groups perform better than the best individuals on letters-to-numbers problems: effects of group size. *Journal of Personality and Social Psychology* 90(4), pp. 644–651.

Lorenz, J., Rauhut, H., Schweitzer, F., Helbing, D. (2011). How social influence can undermine the wisdom of crowd effect. *PNAS*, vol. 108 no. 22, pp. 9020–9025.

Lovett, I. (2013). In California, a Champion for Police Cameras. *New York Times*. August 21, 2013.

Luskin, R.C, Fishkin, J.S, Jowell, R. (2002). Considered Opinions: Deliberative Polling in Britain. *British Journal of Political Science*. Volume 32, issue 03, July 2002, pp. 455-487.

McGonigal, J. (2008). Why I Love Bees: A Case Study in Collective Intelligence Gaming. *The Ecology of Games: Connecting Youth, Games, and Learning*. Edited by Katie Salen. The John D. and Catherine T. MacArthur Foundation Series on Digital Media and Learning. Cambridge, MA: The MIT Press, 2008. 199–228. doi: 10.1162/dmal.9780262693646.199

Mellers, B., Ungar, L., Baron, J., Ramos, J., Gurcay, B., Fincher, K., Scott, S.E., Moore, D., Atanasov, P., Swift, S.A., Murray, T., Stone, E., Tetlock, P.E. (2014). Psychological Strategies for Winning a Geopolitical Forecasting Tournament. *Psychological Science*. DOI: 10.1177/0956797614524255 pp. 1–10.

Moore, M., Hartley, J. (2008). Innovations in governance. *Public Management Review*, 10:1, 3-20, DOI: 10.1080/14719030701763161

Morita, P.P., Burns, C.M. (2013). Designing for Interpersonal Trust – The Power of Trust Tokens. *Proceedings of the human factors and ergonomics society 57th annual meeting - 2013*. 339-343.

Moshman, D., Geil, M. (1998). Collaborative Reasoning: Evidence for Collective Rationality, *Thinking & Reasoning*, 4:3, 231-248

Oxford Martin School. (2013). *Now for the Long Term: The Report of the Oxford Martin Commission for Future Generations*. Oxford Martin School, Oxford University.

Sandia National Laboratories. (2007) "Complex 'Wicked' Problems Better Solved Individually Than Through Internet Groups." *ScienceDaily*, 5 December 2007. Retrieved 2014-05-07. www.sciencedaily.com/releases/2007/11/071130172937.htm

Simon, J., Anna, D. (2013) People Powered Social Innovation: The Need for Citizen Engagement. (2013). *Social Space*. 38-43. Available at: http://ink.library.smu.edu.sg/lien_research/118

Simpson, E., Roberts, S., Psorakis, I., Smith, A. (2013). Dynamic Bayesian combination of multiple imperfect classifiers. *Decision Making and Imperfection*. Studies in Computational Intelligence Volume 474, 2013, pp 1-35

Smith, C.M., Bushouse, E., Lord, J. (2009). Individual and group performance on insight problems: The effects of experimentally induced fixation. *Group Processes & Intergroup Relations* 13(1) 91–99

Stevens, S.M., Dornburg, C.C., Hendrickson, S.M.L., Davidson, G.S. (2008). Individual and Group Electronic Brainstorming In an Industrial Setting. *Proceedings of the human factors and ergonomics society 52nd Annual meeting 2008*. pp. 493-497

Sunstein, C.R. (2006). *Infotopia: How Many Minds Produce Knowledge*. Oxford University Press.

Swallow, E. (2013). The Rise Of The Reputation Economy. *Forbes* 9 October 2013.

Vbuterin. 2014. *A Next-Generation Smart Contract and Decentralized Application Platform*. Retrieved 2014-05-07. <https://github.com/ethereum/wiki/wiki/%5BEnglish%5D-White-Paper>

Williamson, O.E. (1967). Hierarchical Control and Optimum Firm Size. *Journal of Political Economy* Vol. 75, No. 2, Apr., 1967 123-138.

World Economic Forum 2014. *Global Risks 2014*. Ninth Edition. www.weforum.org/risks